# Cray and AMD Scientific Libraries

Mary Beth Hribar, Cray Inc.

Chip Freitag, AMD

CUG 2006

# Libraries for Cray Systems

- Cray LibSci (Mary Beth)
  - XT3 and beyond
  - BlackWidow
  - Current projects
- AMD Core Math Library (Chip)

# Cray Contributors

- Mary Beth Hribar: Manager
- Adrian Tate: Sca/LAPACK, Sparse linear algebra
- Bracy Elton: FFTs
- Chao Yang: BLAS, LAPACK, Direct sparse solvers
- John Lewis: Sca/LAPACK, Sparse linear algebra
- Neal Gaarder: libm
- Catherine Knutson:  Library builds and integration

# Library Goals

- Extend performance and functionality for scalar systems
- Tune vector libraries for BlackWidow architecture
- Collaborate with
  - Researchers to obtain newest algorithms
  - AMD to provide best Opteron libraries

# The Cray Roadmap

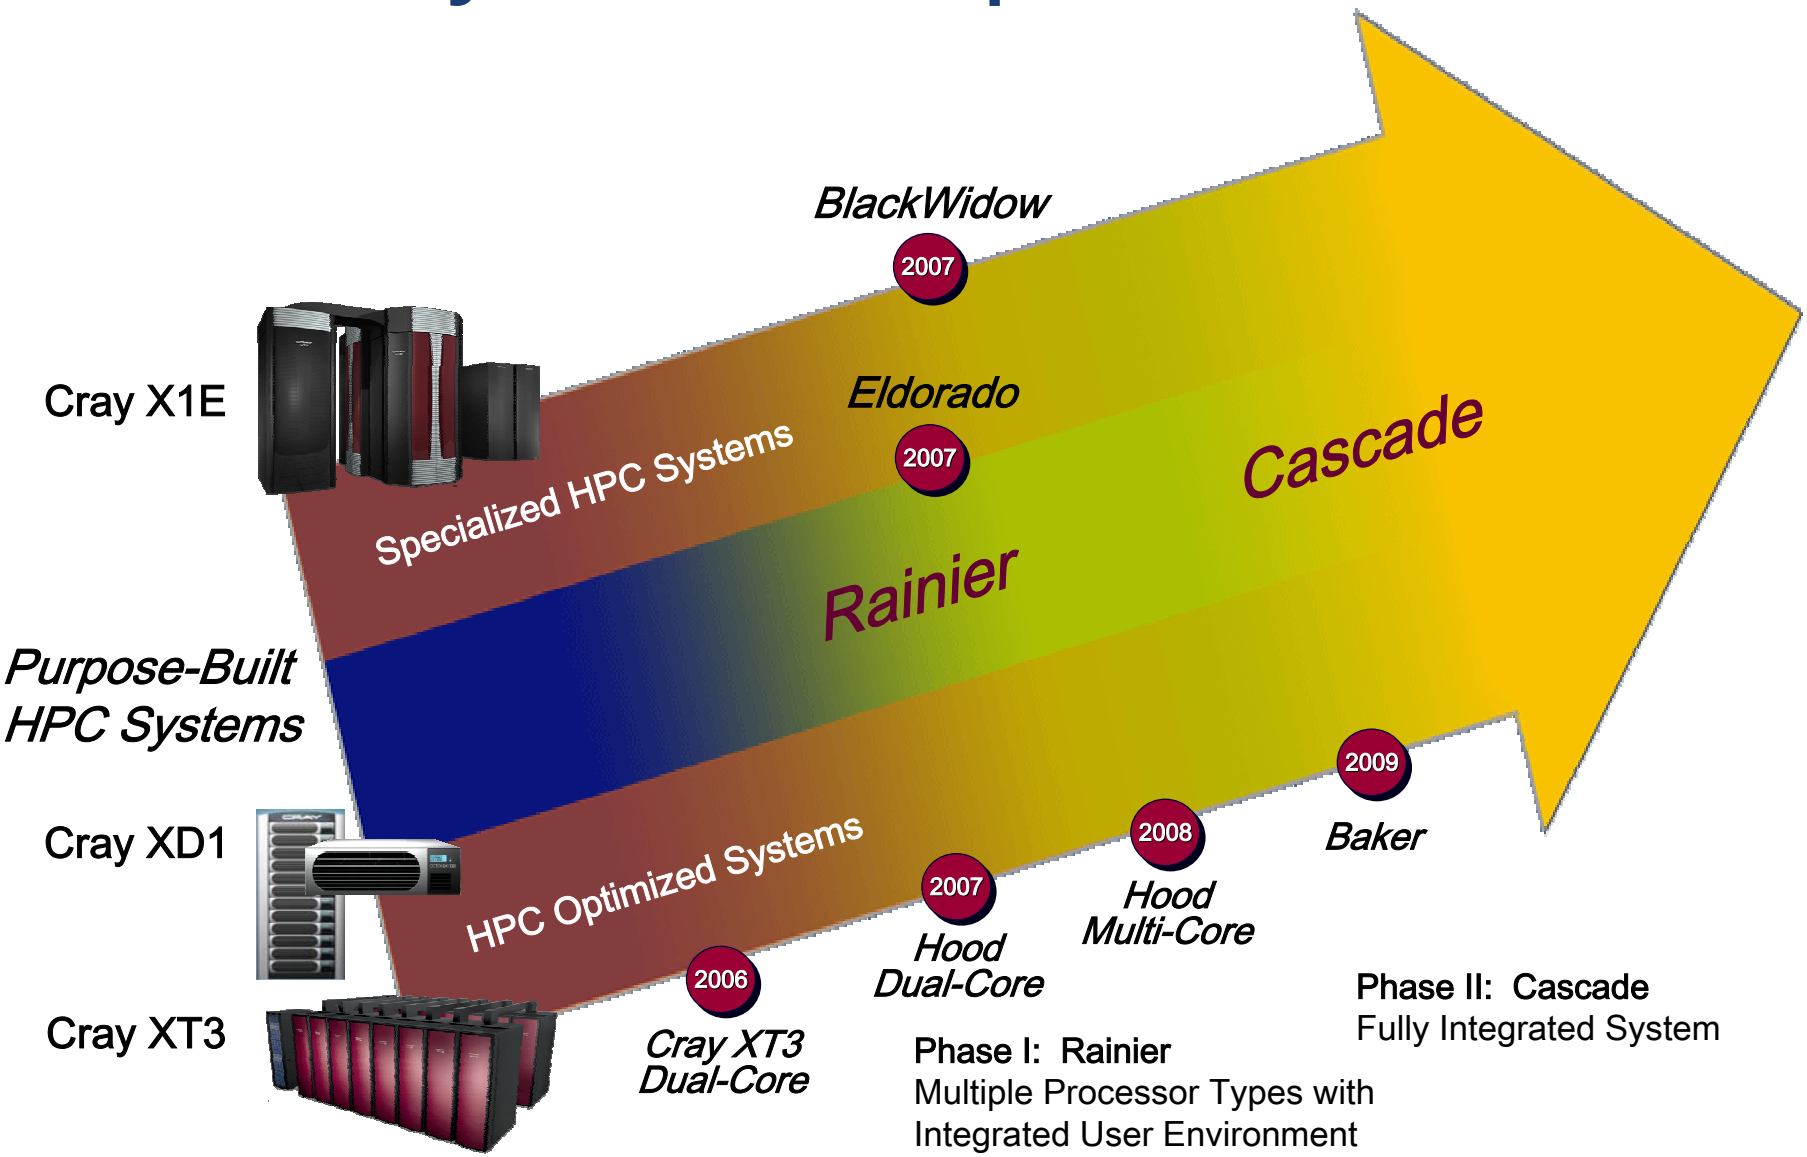


This Presentation May Contain Some Preliminary Information, Subject To Change

# Cray XT3 Libraries

- Cray XT LibSci
  - ScaLAPACK
  - SuperLU
  - Cray FFTs *
- AMD Core Math Library (ACML)
  - BLAS
  - LAPACK
  - ACML FFTs
  - Random number generators
- Goto BLAS
- FFTW 3.1.1

Available in software release 1.5

* Subset of Cray FFTs that map to ACML FFTs

# FFTs on Cray XT3

- FFTs in ACML
  - Provide “plans”
  - Contain OpenMP version
- Add FFTW 3.1.1 in 1.5
- Add Cray FFT interfaces to ACML FFTs in 1.5
- Pre-built plans for FFTW 3.1.1 by end of 2006
- Additional FFT optimizations in 2007

# FFTW on Cray XT3

- Have license to distribute 3.x and 2.1.5 versions
- FFTW 3.1.1
  - Initial release in 1.5
  - Pre-built plans (“Wisdom”) available end of 2006
- FFTW 2.1.5
  - Initial release at end of 2006
  - Included only for MPI FFTW
  - Tuned by demand
- Further optimizations of FFTW 3.x in 2007

# Cray XT3 Sparse Support

- Provide tuned sparse BLAS routines for sparse iterative solvers (end of 2006)
  - PETSc
  - Trilinos
  - User-defined
- Direct sparse solvers
  - SuperLU
  - Pardiso in ACML in 2007

# Libraries for Cray Systems

- Cray LibSci
  - XT3 and beyond
  - BlackWidow
  - Current projects
- AMD Core Math Library

# BlackWidow LibSci

- Extension of X1/X1E LibSci
  - Additional OpenMP support
  - Sparse BLAS
- Further optimizations
  - One-sided communication
  - BlackWidow memory model

# X1/X1E LibSci

**X1/X1E LibSci**

**Single CPU**
**BLAS**
**LAPACK**
**Cray FFTs**
**Cray Direct Sparse Solver**

**SM Parallel (OpenMP)**
**Level 3 BLAS**

**DM Parallel**
**BLACS**
**PBLAS**
**ScaLAPACK**
**Cray DMP FFTs**

Provide four libraries to support streaming and data size options

**MSP/32-bit**
**default**

**MSP/64 bit**
**-sdefault64 (-lsci64)**

**SSP/32-bit**
**-hssp**

**SSP/64-bit**
**-hssp –sdefault64**

# BlackWidow Node

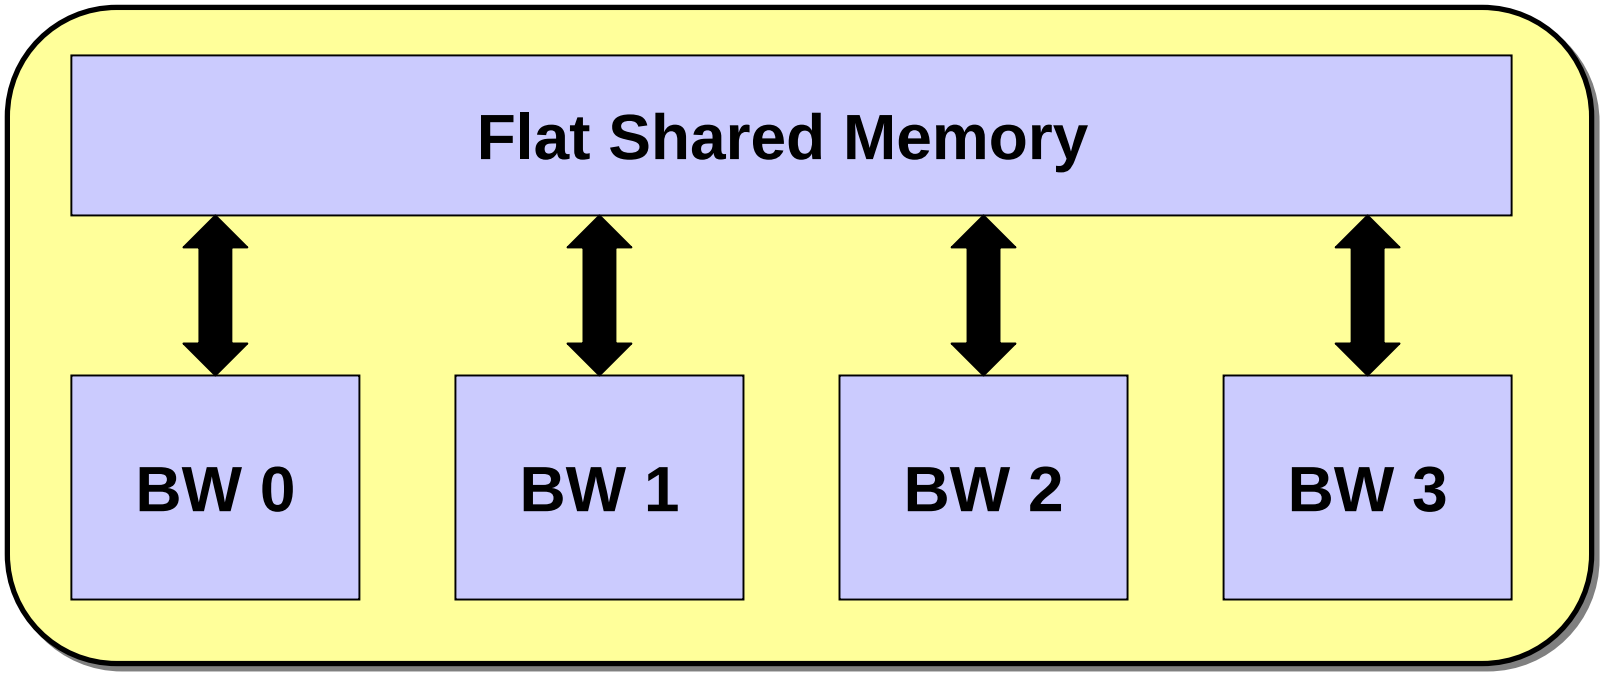


- 4 way SMP
- 2 L1 caches and L2 cache on each processor
- Shared L3 cache
- Faster processing speeds
- Lower memory latency

# BW LibSci

**BW LibSci**

**Single CPU**
BLAS
LAPACK
Cray FFTs
Cray Direct Sparse Solver
Sparse BLAS

**SM Parallel (OpenMP)**
Level 3 BLAS
Cray FFTs
Cray Direct Sparse Solver
LAPACK (subset)

**DM Parallel**
BLACS
PBLAS
ScaLAPACK
Cray DMP FFTs

Provide two libraries to support data size options

**32-bit**
**default**

**64-bit**
**-sdefault64 (-lsci64)**

**New features: Sparse BLAS, OpenMP FFTs/Sparse Solver/LAPACK**

# FFTs

- Cache Tuned
- Added OpenMP capability
- DM FFTs are hybrid
  - Tuned with one-sided communication across nodes
    - CAF and/or SHMEM
  - OpenMP within a node

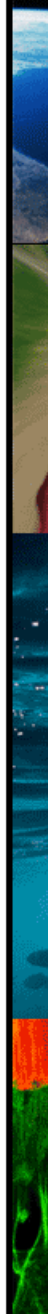

# Sca/LAPACK

- OpenMP versions of LAPACK routines:
  - LU
  - Symmetric Tridiagonalization
  - Cholesky
  - QR
- Collaboration with LBNL to vectorize new eigensolver
- Co-array Fortran PBLAS

# Sparse Computation Support

- Sparse BLAS routines to support iterative solvers in
  - PETSc
  - Trilinos
  - User-defined
- OpenMP direct sparse solvers

# Libraries for Cray Systems

- Cray LibSci
  - XT3 and beyond
  - BlackWidow
  - Current projects
- AMD Core Math Library

# Projects across all Cray systems

- Eigensolver assessment
- ScaLAPACK optimization
- Sparse BLAS strategy

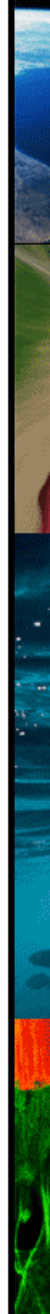

# Eigensolver Support

- Comparison of Sca/LAPACK eigensolvers
  - Divide and Conquer
  - Bisection
  - MRRR ( “Holy Grail”)
  - QR
- Provide Beta release of MRRR for XT3 late 2006
- Work done by Adrian Tate
  - Talk:  Thursday at 9:25 am

# XT3 Matrix 3 10% of spectrum

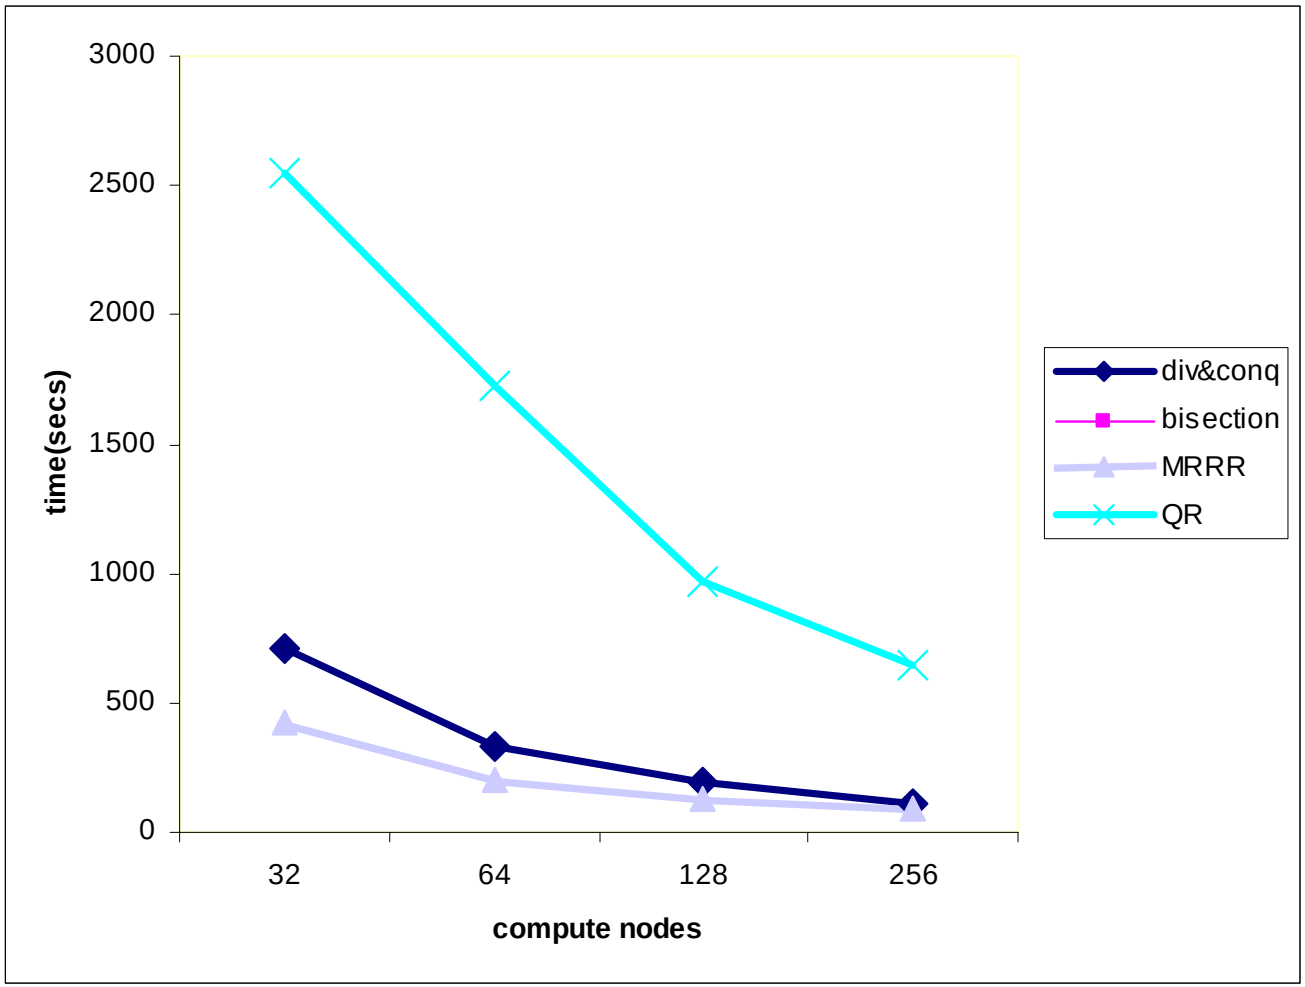

# ScaLAPACK Optimizations

- Exploit fast one-sided communication in
  - BlackWidow
  - Baker
- Introduce more flexible block sizes
  - De-couple two types of blocking
    - Linear algebra block size
    - Distribution block size
  - User can tune two block sizes separately

# Sparse BLAS Routines

- Sparse BLAS (single processor) to support iterative solvers :
  - PETSc
  - Trilinos
  - User-defined
- Investigating Epetra in Trilinos
  - Provides parallel sparse BLAS using
    - Single processor sparse BLAS
    - MPI
  - **Able to tune communication**
  - Possible interface to PETSc
    - One common parallel sparse BLAS implementation
    - Improve communication performance of PETSc solvers

# Library Goals

- Extend performance and functionality for XT3
  - FFTW, Cray FFTs, Goto BLAS
  - Tuned sparse BLAS for better solver performance
- Tune vector libraries for BlackWidow architecture
  - More OpenMP support in LibSci
  - CAF PBLAS
  - Tuned sparse BLAS for better solver performance
- Collaborate with
  - Researchers to obtain newest algorithms
    - MRRR eigensolver
  - AMD to provide best Opteron libraries

# Libraries for Cray Systems

- Cray LibSci
  - XT3 and beyond
  - BlackWidow
  - Current projects
- AMD Core Math Library